



HAL
open science

Estimation of Illuminants From Projections on the Planckian Locus

Baptiste Mazin, Julie Delon, Yann Gousseau

► **To cite this version:**

Baptiste Mazin, Julie Delon, Yann Gousseau. Estimation of Illuminants From Projections on the Planckian Locus. 2013. hal-00915853v1

HAL Id: hal-00915853

<https://hal.science/hal-00915853v1>

Preprint submitted on 9 Dec 2013 (v1), last revised 9 Oct 2014 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Estimation of Illuminants From Projections on the Planckian Locus

Baptiste Mazin, Julie Delon and Yann Gousseau

Abstract—This paper introduces a new approach for the automatic estimation of illuminants in a digital color image. The method relies on two assumptions. First, the image is supposed to contain at least a small set of achromatic pixels. The second assumption is physical and concerns the set of possible illuminants, assumed to be well approximated by black body radiators. The proposed scheme is based on a projection of selected pixels on the Planckian locus in a well chosen chromaticity space, followed by a voting procedure yielding the estimation of the illuminant. This approach is very simple and learning-free. The voting procedure can be extended for the detection of multiple illuminants when necessary. Experiments on various databases show that the performances of this approach are similar to those of the best learning-based state of the art algorithms.

I. INTRODUCTION

The chromaticities of objects in a scene are highly dependent on the light sources. Humans have the ability, known as color constancy, to perceive objects almost independently of the illuminants. In contrast, the influence of illuminants may be a limiting factor for computer vision applications such as object recognition or categorization. The goal of automatic color constancy is to reduce these color variations as much as possible, and to render colors regardless of the light source, mimicking in a way the human capacity. The ability to produce image representations as independent as possible of the illuminants is also necessary for computational photography, where this process is often referred to as *white balance*. The first step of automatic color constancy consists, for a given scene, in estimating the colors of the illuminants. In a second step, when dealing with a single illuminant, the image can be globally corrected using this estimate so that it appears as if taken under a canonical illuminant. In this paper, we focus on the illuminant estimation step. We take interest in situations where one (single illuminant case) or several (multi-illuminant case) light sources are involved in the scene.

A. Single illuminant estimation

Illuminant estimation has been a very active research field in the past decades, see for instance the recent survey [25]. In the case of a single light source, the simpler methods rely only on the distribution of colors in images. For example, a very popular way to estimate illuminant chromaticities in images is to assume that under a canonic light, the average RGB value observed in a scene is grey. This assumption gives rise to the Grey-World algorithm [7], which consists in computing the average color in the image and compensating for the deviation due to the illuminant. Several refinements have been proposed,

restraining the hypothesis to well chosen surfaces of the scene, that are assumed to be grey [39]. Potentially grey pixels can also be selected as the points close to the Planckian locus¹ in a given chromaticity space [37]. This approach may fail when two potentially grey surfaces are present in the scene.

Alternatively, the White-Patch assumption [38] supposes that a surface with perfect reflectance is present in the scene and that this surface correspond to the brighter points in the image. This results in the well-known Max-RGB algorithm, which infers the illuminant by computing separately the maxima of the RGB channels. Variants propose to perform a preliminary filtering of the image [1]. A unified framework has been proposed for these two assumptions (Grey-World and White-Patch) in [22]. More recently, the same unified framework has been extended to the reflectance first and second order derivatives in the so-called Grey-Edge methods [48], [28], [8].

Another way to estimate the chromaticity of light sources is to rely on the dichromatic reflectance model [14], which takes into account the specular part of the reflectance at each pixel. Under this model, the pixels of a monochromatic surface belong to a plane in the RGB color space of the camera, or to a line in a well chosen chromaticity space, see [21]. Finding the intersection of those planes permits to recover the illuminant. In the case of an image containing only one monochromatic object, the illuminant can still be recovered by assuming that the light source chromaticity belongs to the Planckian locus.

Physical constraints on the set of illuminants and on the scene content are also used by Sapiro in [45], where it is assumed that reflectance spectra and illuminant spectral power distributions can be approximated by linear combinations of a small number of known basis functions. A voting procedure is then proposed, where each pixel votes for a set of possible illuminants. The Hough transform permits to retrieve the most probable illuminant.

More involved techniques rely on a learning phase, where color statistics are learned on a training set of images, taken under known light sources. Among these methods, gamut-based approaches rely on the idea that the range of observable colors under a given illuminant is limited [23]. These approaches require to learn the set of observable colors under a known illuminant and to find the best feasible mapping between the colors of the observed scene and this set. The chromaticity of the unknown illuminant is then deduced from this mapping. The gamut mapping problem remains nonethe-

¹In a chromaticity space, the *Planckian locus* refers to the set of black body radiators chromaticities.

less ill-posed, and additional hypotheses are generally necessary [16], [25]. For instance, these hypotheses may concern the diversity of colors in the observed scene [23], or the set of possible light sources [16]. In order to go further, some authors explore not only the range but also the distribution of colors under canonical illuminants [20], [43]. This permits to recast the illuminant estimation as a maximum likelihood problem [20]. Other methods also involve a simple learning of color statistics [24].

While they offer correct estimations in a large number of situations, all of the previously mentioned approaches present different failure cases. A last category of illuminant estimation techniques thus proposes to combine the output of several of these methods [49], [4], [30]. The optimal combination strategy is learned on training image datasets, eventually relying on the statistical properties of natural images, or on more involved semantic information. Recently, Vasquez-Corral *et al.* [50] propose to refine color constancy approaches by incorporating perceptual constraints, weighting the set of illuminants “according to their ability to map the corrected image onto specific colors”, chosen as universal color categories.

B. Multi illuminant estimation

All of the previous methods assume that only one illuminant is present in the scene and that this illuminant is uniform. Less attention has been paid in the literature to the multi-illuminant case, or to the case of a varying illuminant, although they can be considered as more realistic in practice. Typical examples are the case of outdoor images, with objects illuminated by direct sunlight and objects in shadows receiving light scattered by the sky, or the case of indoor scenes with both an incandescent light source and a natural light source coming from a window.

The Retinex theory proposed by Land and McCann [38] is one of the first model addressing the problem of local color constancy. In order not to confuse changes in illumination from changes in reflectance, the method assumes that the illuminant varies spatially smoothly: small changes are thus interpreted as illumination variations, while large ones are supposed to be due to reflectance changes. Several papers draw on the same ideas and try to incorporate additional constraints to the problem, either on the set of possible illuminants [18], [2] (assumed to be on the Planckian locus) or on the set of surfaces reflectances [2]. These approaches [18], [2] also exploit the illumination variations along matte surfaces to constrain the solution². The work of Finlayson [18] is extended in [36], by including noise in the model, and refined in [35] for outdoor scenes, relying on the differences between shadowed and non-shadowed regions.

In [13], Ebner proposes a very simple solution to the local illuminant estimation, using a local and iterative version of the grey world algorithm to compute an illuminant estimation at each pixel. This algorithm is refined in [15] by using a non-uniform averaging of local pixels to address sudden illuminant variations. More recently, the authors of [5] propose to apply

²These surfaces must first be identified, for instance thanks to a segmentation step [2].

different color constancy algorithms on super-pixels and to combine their results. In a similar way, Gijensij *et al.* [29] evaluate color constancy algorithm on image patches, sampled with different strategies.

Some authors uses specular highlights [33] to estimate the spectral power distribution of illuminants through a spectral imaging device.

Finally, different user guided methods have been proposed in the literature [32], [6]. The method described in [32] permits to estimate for each image pixel the mixture between two illuminants specified by the user. This approach works well when the transitions between illuminants are smooth but fails in case of spatially abrupt illuminant changes. To address this problem, Boyadzhiev [6] proposes an algorithm allowing to correct images based on user indications on the nature of the surfaces, specifying whether they are neutral or uniform.

C. Contributions of the paper

The goal of this paper is to introduce a simple and effective method for the estimation of light sources chromaticities, without relying on any learning phase. Our approach draws on a novel voting procedure, combined with two physical assumptions, one concerning the content of the observed scene and the other concerning the set of possible illuminants. First, it is assumed that the scene contains at least some achromatic (or perfectly reflective) pixels; second, the feasible illuminants chromaticities are supposed to be in the vicinity of the Planckian locus. A refined version of the proposed voting procedure permits to deal with situations where more than one illuminant is involved in the scene, and to automatically estimate the number of these illuminants. The whole algorithm does not require any learning step and uses intuitive parameters. The method being applied globally on the image, the result varies only slightly with the scales of images. Various experiments show the efficiency of the proposed method when compared to state of the art approaches on different databases. A conference proceedings version of the first part of this work was published in [40].

The paper is organized as follows. In Section II, we recall the color formation model and discuss color constancy. Section III describes the main steps of our color constancy algorithm in the case where only one illuminant has to be estimated. The multi-illuminant version is detailed in Section IV. In Sections V-B and V-C, the ability of the algorithm to retrieve light sources chromaticities is tested on several databases, both in the mono-illuminant and multi-illuminant cases.

II. ILLUMINANTS AND COLORS IN IMAGES

This section presents a brief reminder on color image formation and on color constancy. We also recall the concept of black-body radiator, which is used in the rest of the paper to constrain the set of possible illuminants.

A. Color image formation

Let us denote by $E(\cdot, \mathbf{m})$ the irradiance spectrum falling on the camera sensor at pixel \mathbf{m} and by $\rho(\lambda) =$

$(\rho_R(\lambda), \rho_G(\lambda), \rho_B(\lambda))$ the camera sensitivities functions. The vector $\mathbf{p}(\mathbf{m}) = (p_R(\mathbf{m}), p_G(\mathbf{m}), p_B(\mathbf{m}))$ measured by the camera at \mathbf{m} is given by

$$p_c(\mathbf{m}) = \int_{\Lambda} E(\lambda, \mathbf{m}) \rho_c(\lambda) d\lambda, \quad c \in \{R, G, B\} \quad (1)$$

where Λ is the visible spectrum. The different image formation models differ in the way they describe the spectrum $E(\cdot, \mathbf{m})$ reaching the camera sensor at pixel \mathbf{m} . One of the most popular descriptions is given by the dichromatic reflection model [46], which decomposes $E(\cdot, \mathbf{m})$ in two parts, a Lambertian component and a specular component:

$$E(\lambda, \mathbf{m}) = L(\lambda, \mathbf{m})(\alpha_l(\mathbf{m})S(\lambda, \mathbf{m}) + \alpha_s(\mathbf{m})), \quad (2)$$

where $L(\cdot, \mathbf{m})$ is the illuminant spectrum at \mathbf{m} , $S(\cdot, \mathbf{m})$ is the object reflectance spectrum at \mathbf{m} and where $\alpha_l(\mathbf{m})$ and $\alpha_s(\mathbf{m})$ weight the Lambertian and specular components. The simpler Lambertian model ignores the specular term.

The illuminant estimation problem is intrinsically under-constrained: even in the case of a single illuminant, uniform across the scene, it is not possible to recover the complete power distribution $L(\lambda)$ from the discrete sensor measures $p_c(\mathbf{m})$. In practice, color constancy only aims at computing for each given pixel \mathbf{m} the triplet of values $(\int_{\Lambda} \rho_c(\lambda) L(\lambda, \mathbf{m}) d\lambda)_{c \in \{R, G, B\}}$. Estimating these values from the observations is still an ill-posed problem and generally necessitates strong assumptions on the image content or on the illuminant properties, as described in the following paragraphs.

B. Color correction and the Von Kries diagonal model

In practice, it is usual to assume that the illuminant spectrum L is constant over the whole image and that the camera sensitivity functions can be approximated by Dirac delta functions³ associated to wavelengths λ_R, λ_G and λ_B . Under these hypotheses, Equation (1) can be rewritten as

$$p_c(\mathbf{m}) = L(\lambda_c)(\alpha_l(\mathbf{m})S(\lambda_c, \mathbf{m}) + \alpha_s(\mathbf{m})), \quad c \in \{R, G, B\}. \quad (3)$$

In this case, the illuminant estimation problem amounts to retrieve $(L_R, L_G, L_B) := (L(\lambda_R), L(\lambda_G), L(\lambda_B))$. Assuming that this illuminant has been retrieved, computing a version of the same image taken under another illuminant L^0 amounts to multiply the vector $(p_R(\mathbf{m}), p_G(\mathbf{m}), p_B(\mathbf{m}))$ by the diagonal matrix [42]

$$\begin{pmatrix} \frac{L_R^0}{L_R} & 0 & 0 \\ 0 & \frac{L_G^0}{L_G} & 0 \\ 0 & 0 & \frac{L_B^0}{L_B} \end{pmatrix}. \quad (4)$$

C. Distribution of feasible illuminants

It is often necessary to impose a priori information on the chromaticity of feasible illuminants to solve the color constancy problem. A possible hypothesis in this context is to assume that the possible illuminants are well represented

³This assumption, very common in the literature, is only an approximation. Nonetheless, observe that a sharpening technique can be applied if the sensors response functions are not sufficiently narrow-band [17].

by black-body radiators. According to the Planck model, the spectrum of light emitted by such an idealized physical body only depends on its temperature T and is given by

$$L(T, \lambda) = c_1 \lambda^{-5} \left[\exp\left(\frac{c_2}{\lambda T}\right) - 1 \right]^{-1}, \quad (5)$$

where T is the temperature in kelvins, λ is the wavelength and $c_1 = 3.74183 \times 10^{16} \text{ Wm}^2$ and $c_2 = 1.4388 \times 10^{-2} \text{ mK}$ are two constants.

Planck's formula is an accurate estimation of the spectrum of most natural illuminants including daylight [34]. Besides, even if artificial light sources (such as fluorescent lights) generally have a power spectrum which is not well represented by the Planck formula, it was shown that such sources produce a set of perceived colors which is very close to those produced by a black-body radiator. Indeed, as underlined in [19], "for almost all daylights and typical man-made lights, including fluorescents, there exists a black-body radiator (...) which (...) will induce very similar RGBs for most surface reflectances".

D. Chromaticity spaces

When solving the color constancy problem, it is not possible to distinguish which part of the observed intensity comes from the light source and which part comes from the object reflectance. For this reason, the intensity information at \mathbf{m} cannot be recovered and the illuminant (L_R, L_G, L_B) is only estimated up to a multiplicative factor. When trying to retrieve this illuminant, it is thus sensible to work in a chromaticity space, *i.e.* a 2D representation of the 3D color space discarding the intensity component. For instance, the *rg* chromaticity is obtained by normalizing the *RGB* triplet into $\left(\frac{R}{R+G+B}, \frac{G}{R+G+B}\right)$. The *xy* chromaticity space is obtained in the same way from the CIE 1931 *XYZ* color space. Figure 1 shows the curve of black body radiators chromaticities, called the Planckian locus, in the CIE 1931 *xy* chromaticity diagram (*i.e.* the set of visible colors in the *xy* space). We also display on the same diagram the chromaticities of a large variety of illuminants (obtained from the Colorchecker database, described in [24]). Observe that all of these chromaticities lie on (or close to) the Planckian locus.

In this paper, we make use of the CIE 1960 *uv* diagram, which was proposed by MacAdam to define correlated color temperatures (CCT)⁴. According to the CIE, the isotherms (lines of constant CCT) are defined in this space as the lines perpendicular to the Planckian locus [51]. It follows that the correlated color temperature $T(\mathbf{m})$ of a pixel \mathbf{m} is the one of its orthogonal projection on the Planckian locus, *i.e.* the one of its nearest black body radiator.

III. VOTING FOR THE MOST PROBABLE ILLUMINANT

In this section, we present a first version of our illuminant estimation algorithm. We assume in the following that a single illuminant is used. As described above, estimating the triplet (L_R, L_G, L_B) from the measures $p_c(\mathbf{m})$ is an underconstrained problem and it necessitates some additional

⁴The correlated color temperature of a light source is the temperature of the black body radiator with the most perceptually similar color.

assumptions. We rely here on two hypotheses. The first one concerns the set of possible illuminants of a scene. We assume that this set is not far from the Planckian locus, which means that the illuminants are supposed to be well approximated by black-body radiators. This *a priori* is quite common in the literature and was shown to be a sane assumption for most illuminants [34], [19]. Our second assumption concerns the scene content and is inspired by the white-patch method: we assume that the scene contains some achromatic or perfectly reflective surfaces. This assumption has, in our opinion, a stronger physical justification than the one of the grey-world algorithm (the average reflectance of a scene is grey). While, of course it is possible to find or create counter examples, we observed that this hypothesis is very often satisfied in practice, especially for scenes containing man-made objects. Both hypotheses will be discussed in the section devoted to the algorithm evaluation.

The main idea of our method is to select pixels close to the Planckian locus and to make these pixels vote for the most probable illuminant. Before describing the algorithm more precisely, an important detail must be specified on the color space used for the input images. The RGB space of a camera depends heavily on the sensitivity functions of the sensor. While it is possible to estimate the illuminant directly in these camera-specific color spaces, it becomes difficult to assess the quality of the results, as error measures do not have the same meaning between two different RGB spaces. In order to obtain a consistent algorithm between cameras, we believe it is important that the input images are provided in a universal space as sRGB (observe that most camera manufacturers allow the user to save images in this space).

The first part of the algorithm works as follows. First, for all pixels \mathbf{m} , the observed color triplet $p_c(\mathbf{m})$ ($c \in \{R, G, B\}$) is transformed in a 2D chromaticity vector $\mathbf{c}(\mathbf{m}) = (u(\mathbf{m}), v(\mathbf{m}))$ in the CIE 1960 uv diagram. As explained before, working with chromaticities is sensible since the intensity information at \mathbf{m} cannot be recovered. Next, all pixels at a distance from the Planckian locus \mathcal{L} larger than a threshold δ are discarded, which means that we only keep pixels \mathbf{m} such that

$$\min_{\mathbf{r} \in \mathcal{L}} \|\mathbf{c}(\mathbf{m}) - \mathbf{r}\|_2 < \delta, \quad (6)$$

these pixels being considered as the most probably grey in the image. We then project each of the remaining chromaticities $\mathbf{c}(\mathbf{m})$ orthogonally on the Planckian locus:

$$\mathbf{c}_{Planck}(\mathbf{m}) := \text{proj}_{\mathcal{L}}(\mathbf{c}(\mathbf{m})) = \arg \min_{\mathbf{r} \in \mathcal{L}} \|\mathbf{c}(\mathbf{m}) - \mathbf{r}\|_2. \quad (7)$$

If the argmin in Equation (7) is not uniquely defined, we choose randomly one of the possible projections. This assignment permits to define the correlated color temperature $T(\mathbf{m})$ of the pixel \mathbf{m} , as the temperature of the black body radiator whose chromaticity is $\mathbf{c}_{Planck}(\mathbf{m})$. We keep only the pixels whose correlated temperature falls inside a given interval $[T_{min}, T_{max}]$, containing the temperatures of the most common illuminants.

The second step of the algorithm is a voting procedure permitting to retrieve the most probable temperature T^0 for the

illuminant. At the end of the first step, the remaining points in the chromaticity diagram are used to build a temperature histogram. For computing this discrete histogram, we convert the correlated temperature T (measured in Kelvin) into the MIREd (Micro Reciprocal Degree) scale $T_{Mired} = 10^6/T$ before the discretization. Indeed, the Kelvin temperature scale does not yield a satisfying sampling of the Planckian locus. The conversion to the MIREd scale is advocated by Priest [41] as a good way to sample color temperature according to human perception. When computing the temperature histogram, the contribution of each pixel is weighted by a power of its luminance. This weighting is a way to favor the brightest pixels in the image, and to limit the influence of the darkest ones without completely discarding them⁵. The temperature T^0 of the illuminant is then estimated as the most represented one in this weighted histogram. For the sake of simplicity, the argmax of the histogram is chosen here, but a more sophisticated procedure, based on mode detection, can also be applied. We will see in the next Section how such a procedure can be used for the detection and estimation of several illuminants. At the last step of the algorithm, we compute the barycenter of all pixels \mathbf{m} who have participated in the choice of T^0 (those pixels which were kept to build the histogram and whose discrete correlated color temperature falls in the bin T^0). This barycenter, which may be outside of the Planckian locus but remains close to it, is chosen as the output illuminant. Algorithm 1 gives a complete description of the procedure.

IV. ALGORITHM REFINEMENT WITH MODES SELECTION

The algorithm described in the previous section outputs a single illuminant after selecting the principal mode of the temperature histogram in a rough way (argmax of the histogram). In practice, the histogram may contain multiple modes, especially when several illuminants are present in the scene. Moreover, even with a single illuminant, the argmax choice is not always the most appropriate, since the mode may extend over a large range of temperatures. In this section, we propose a mode selection algorithm, inspired from *a contrario* methods [12], which permits to refine the detection and to manage the multi-illuminant cases.

Since we do not have any *a priori* information on the number of illuminants present in the scene, segmentation-based approaches such as k-Means cannot be used straightforwardly. Gaussian mixtures with a minimum description length criteria can be used to address the problem, but this would introduce an *a priori* on the shape of the modes that is not very accurate in practice. We propose here a generic algorithm which detect modes in an histogram. The algorithm is inspired from the *a contrario* mode selection presented in [12], [11]. It is adapted to the specific weighting procedure used to build the temperature histogram.

Let H be an histogram with N bins, each bin corresponding to a CCT interval in the mired scale. This histogram is obtained as explained in Algorithm 1: the contribution of each pixel \mathbf{m}_i $i \in \{1, \dots, M\}$ is weighted by a mass w_i equal to a power

⁵In a way, this weighting scheme can be seen as a trade off between the Grey-World and the White-Patch hypotheses [22].

input : Image \mathbf{p} in the space sRGB. Thresholds δ ,
 T_{\min} , T_{\max} , number N of bins in the histogram,
power n , canonical illuminant $(u^{\text{ref}}, v^{\text{ref}})$.
output: Chromaticity (u^e, v^e) estimated in the CIE 1960
 uv space.
Initialize the histogram $H = 0$ on N bins;
for each pixel \mathbf{m} **do**
 $\mathbf{c}(\mathbf{m}) = \text{Convert}_{s\text{RGB} \rightarrow uv}(p_c(\mathbf{m}))$;
 $[\mathbf{c}_{\text{Planck}}(\mathbf{m}), T(\mathbf{m})] = \text{Proj}_{\text{Planck}}(\mathbf{c}(\mathbf{m}))$;
 // $\mathbf{c}_{\text{Planck}}(\mathbf{m})$ is the projection of the
 chromaticity $\mathbf{c}(\mathbf{m})$ on the Planckian
 locus in the uv chromaticity space
 and $T(\mathbf{m})$ is the temperature (in
 Kelvins) of the projection;
 $d(\mathbf{m}) = \text{distance}(\mathbf{c}(\mathbf{m}), \mathbf{c}_{\text{Planck}}(\mathbf{m}))$;
 if $d(\mathbf{m}) < \delta$ **and** $T_{\min} \leq T(\mathbf{m}) \leq T_{\max}$ **then**
 $T_{\text{MIREDD}}(\mathbf{m}) = 10^6 / T(\mathbf{m})$;
 $\text{weight}(\mathbf{m}) = \text{luminance}(\mathbf{m})^n$;
 Add $\text{weight}(\mathbf{m})$ to bin $T_{\text{MIREDD}}(\mathbf{m})$ in H ;
 end
end
if $H == \emptyset$ **then**
 return $(u^{\text{ref}}, v^{\text{ref}})$;
end
Find $T^0 = \arg \max(H)$;
 $P = \text{list of pixels } \mathbf{m} \text{ such that } T_{\text{MIREDD}}(\mathbf{m}) = T^0 \text{ and}$
 $d(\mathbf{m}) < \delta$
return $(u^e, v^e) = \text{Barycenter}(P)$;

Algorithm 1: Illuminant selection algorithm. In practice, the following parameters provide good results on various dataset $\delta = 0.0125$, $T_{\min} = 2000K$, $T_{\max} = 20,000K$, $N = 30$, and $n = 3$.

function of its luminance. Following the general idea of a *contrario* approaches [12], modes are detected as intervals in which the total mass contradicts a null hypothesis \mathcal{H}_0 . Let us describe more precisely this hypothesis \mathcal{H}_0 . For a given interval $[a, b]$ and a pixel \mathbf{m}_i we define two random variables B_i and W_i . The first one B_i , corresponds to the event “the CCT of \mathbf{m}_i belongs to $[a, b]$ ”. The second random variable W_i corresponds to the mass w_i of \mathbf{m}_i . The variables B_i and W_i $i \in \{1, \dots, M\}$ are said to follow the null hypothesis \mathcal{H}_0 if they are mutually independent, if all variables B_i follow the same Bernoulli law of parameter $p_{[a,b]}$ ($p_{[a,b]}$ being the relative length of the interval $[a, b]$ in the histogram) and if all the variables W_i follow an exponential law with expectation λ . The exponential distribution have been chosen because it is a reasonable approximation of the law of weights w_i on most natural images. In practice, the parameter λ is learned for each image. Now, let us consider the random variable $D_{[a,b]} = \sum_{i=1}^M W_i B_i$. This variable describes the mass observed in the interval $[a, b]$ when the null hypothesis is satisfied. When M is large enough, the central limit theorem can be applied and leads to approximate the law of D with a Normal law $\mathcal{N}(Mp_{[a,b]}\lambda, Mp\lambda^2 * (2 - p_{[a,b]}))$. It follows that for all α , we can approximate the probability $\mathbb{P}[D_{[a,b]} > \alpha]$

by

$$F(\alpha) := \int_{\alpha}^{+\infty} g_{\mu, \sigma}(x) dx, \quad (8)$$

with $\mu = Mp_{[a,b]}\lambda$ and $\sigma = Mp\lambda^2 * (2 - p_{[a,b]})$. In consequence, the interval $[a, b]$ is said to be meaningful if it contains a mass $H_{[a,b]}$ such that

$$F[H_{[a,b]}] < \frac{2}{N(N-1)}. \quad (9)$$

The coefficient $\frac{N(N-1)}{2}$, corresponding to the number of intervals in the histogram H , is introduced to control the number of false alarms [12], [11]. Indeed, thresholding at $\frac{2}{N(N-1)}$ permits to ensure that the expected number of detected modes under the null hypothesis is smaller than one.

Intervals satisfying this test are selected. Then, a maximality criterion is applied, permitting to keep only intervals which neither contain, nor are contained in more meaningful intervals (intervals with a smaller $F[H_{[a,b]}]$ value). For each of the maximal intervals, we apply the final procedure described in Algorithm 1 and for each mode selected, we average the contributing pixels to obtain an illuminant estimation.

V. EXPERIMENTS

In this section, we investigate the performance of our illuminant estimation algorithm. Our method is first evaluated and confronted with standard or state-of-the-art approaches on scenes containing a single illuminant. In order to assess the quality of the results, the comparison is performed on three different databases, namely the Colorchecker dataset [24], the Grey-Ball dataset [9] and the SFU laboratory dataset [3]. In the second part of this section, we test the ability of our method to deal with multi-illuminant situations.

A. Colorspaces

Before describing the experimental protocol used in the experiments, it is important to specify which color spaces are used in these experiments and why. As already explained in Section III, each camera has its own RGB color space, which depends on the sensitivity functions of its sensors. As a consequence, RGB measurements of two different cameras are not directly comparable. Their respective color components may have very different dynamic ranges for instance. In order not to depend on these specific color spaces, our algorithm works with images provided in the sRGB color space. This has several distinct advantages. First, the algorithm parameters do not have to be adapted to the camera color space and can be defined once and for all in sRGB. Second, the computation of color distances and the projections on the Planckian locus are also much more sound in universal color spaces than in camera-specific spaces. In the same way, one can easily convert sRGB colors to uv chromaticities, which are recommended [51] by the CIE in order to compute correlated color temperatures⁶. Finally, working with a universal space is the only way to obtain error measures between illuminant

⁶In practice, we use the approximation proposed by [31] to compute the projections and correlated color temperatures.

estimations which are consistent from one camera to the other. We believe that this consistency between images is essential since error measures will be used to compute error statistics for different databases. In practice, when images are provided in their camera-specific color spaces, it is necessary for our algorithm to convert these images to sRGB. Fortunately, most recent cameras can record images in universal color spaces and allow a conversion between their own colorspace and sRGB.

B. Single illuminant estimation

1) *Protocol*: For the single illuminant evaluation, we follow an experimental protocol similar to the one described in [26]. Each image used in these experiments is provided with a measure \mathbf{L}^g of the light source illuminating the scene. Following [27], the angular error between this groundtruth and the estimated illuminant L^e is defined as

$$E_a(\mathbf{L}^e, \mathbf{L}^g) = \cos^{-1} \left(\frac{\mathbf{L}^e \cdot \mathbf{L}^g}{\|\mathbf{L}^e\| \cdot \|\mathbf{L}^g\|} \right). \quad (10)$$

This error can be seen as an angle in a RGB space between two “grey” axes defined by the illuminants \mathbf{L}^e and \mathbf{L}^g . Observe that the illuminant intensity is not involved in this error computation.

We compute several error statistics for each database. As noticed by the authors of [26], the error distribution is not symmetric and therefore, the mean is weakly informative. The median and the trimean measure are more meaningful statistics in practice.

2) *Databases*: We now describe the three databases used for the single illuminant comparison.

a) *Colorchecker dataset*: The Colorchecker dataset, provided by Gehler *et al.* [24], is composed of 568 images, captured with two different cameras (a Canon EOS-1D and a Canon EOS-5D). As a consequence, as said before, results computed in the camera RGB space, averaged in different spaces, should be taken cautiously. The database is composed of indoor and outdoor scenes, corresponding to a large variety of illuminants. These groundtruth illuminants are provided as RGB triplets in the respective colorspace of the cameras. The two cameras used to create the dataset being known, we can easily convert these triplets in the sRGB colorspace⁷ and visualize their projections on the xy chromaticity diagram (see Figure 1).

Gehler *et al.* [24] provide the database in both RAW and TIFF versions. The TIFF version was found [47] to be subject to non-linear corrections such as gamma-correction, clipping and demosaicing. Our experiments therefore rely on the RAW versions, which are converted to sRGB. In order to apply this conversion properly, we first remove the black level of the camera (129 for the Canon 5D and 0 for the Canon 1D).

b) *Grey-Ball dataset*: The Grey-Ball dataset [9] comes with 11,346 JPEG images extracted from videos. As recommended by [26], the gamma correction has been removed in order to work with linear images. The output RGB system of the camera (Sony VX-200) being unknown, we assume a sRGB space for these images. For each image, a grey-ball




Fig. 1. Colorchecker illuminants groundtruth in the xy diagram.

is present in the visual field of the camera. The creators of the database use the pixels belonging to the grey-ball in order to deduce the groundtruth illuminants and to provide these illuminants as RGB triplets⁸. As shown on Figure 2, this dataset covers a large variety of illuminants.

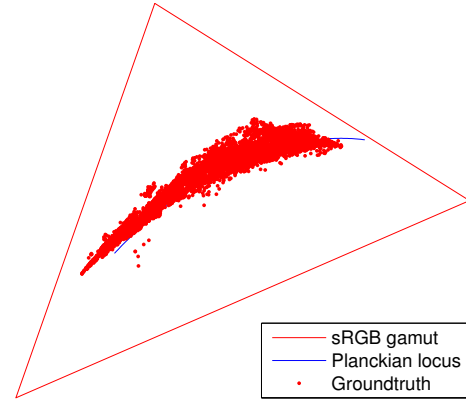


Fig. 2. Grey-Ball groundtruth illuminants projected in the xy chromaticity diagram.

c) *SFU laboratory dataset*: This database [3] contains 321 images from 31 different scenes, each one being illuminated by 11 light sources. Scenes generally contain only few objects, the range of colors in images is thus limited. This database is therefore a challenging case for color constancy algorithms.

For each image in the database, the authors provide the spectra of the groundtruth illuminants and the corresponding triplets in the colorspace of the camera. When projecting these spectra onto the xy diagram, as shown in the left part of Figure 3, we observe that all of these illuminants are in the vicinity the Planckian locus. One image in the database contains a colorchecker. We use this image and the knowledge of the corresponding illuminant to estimate the conversion matrix between the camera colorspace and sRGB⁹. We use

⁸Pixels belonging to the grey-ball are discarded from the comparative experiments.

⁹Note that this conversion matrix could also be estimated thanks to the camera sensitivity functions. However, these sensitivity functions are rarely known in practice, while it is always possible to take a picture of a colorchecker under a known illuminant.

⁷The conversion matrices are documented in dcrw [10].

TABLE I

FIRST TABLE SHOWS PERFORMANCES ON THE COLORCHECKER DATASET [47] IN THE CAMERA RGB SPACE AND SECOND TABLE SHOWS RESULTS IN THE sRGB COLORSPACE. RESULTS ARE BOTH COMPARED TO LEARNING-FREE METHODS ([38], [7], [3], [48]) AND METHODS INVOLVING A TRAINING PHASE ([23], [25], [8]).

Method	Mean	Median	Trimean	Best-25%	Worst-25%
White-Patch [38]	7.5°	5.7°	6.4°	1.5°	16.2°
Grey-World [7]	6.4°	6.3°	6.3°	2.3°	10.6°
general Grey-World [3]	4.7°	3.5°	3.8°	1.0°	10.2°
1st – order Grey-Edge [48]	5.3°	4.5°	4.7°	1.8°	10.2°
Pixel-based Gamut Mapping [23]	4.2°	2.3°	2.9°	0.5°	10.8°
Edge-based Gamut Mapping [25]	6.5°	5.0°	5.4°	1.9°	13.6°
ML (category-wise prior) [8]	3.7°	3.0°	3.1°	1.0°	7.6°
Proposed method $\delta = 0.02$	4.1°	2.7°	3.1°	0.8°	9.6°
Proposed method (fixed parameters)	4.5°	3.1°	3.5°	0.8°	10.8°

Method	Mean	Median	Trimean	Best-25%	Worst-25%
White-Patch [38]	9.6°	7.7°	8.6°	2.2°	20.3°
Grey-World [7]	9.1°	8.6°	8.8°	2.9°	16.2°
general Grey-World [3]	6.6°	4.9°	5.4°	1.4°	14.6°
1st – order Grey-Edge [48]	7.0°	6.3°	6.4°	2.8°	12.5°
Pixel-based Gamut Mapping [23]	5.6°	3.4°	4.0°	0.8°	13.8°
Edge-based Gamut Mapping [25]	9.6°	7.4°	8.1°	2.8°	20.0°
ML (category-wise prior) [8]	5.0°	4.1°	4.3°	1.4°	10.1°
Proposed method $\delta = 0.02$	5.2°	4.0°	4.3°	1.4°	10.8°
Proposed method (fixed parameters)	6.1°	4.2°	4.8°	1.2°	14.1°

TABLE II

PERFORMANCES ON THE LINEAR SFU GREY-BALL DATABASE, IN THE CAMERA COLORSPACE.

Method	Mean	Median	Trimean	Best-25%	Worst-25%
White-Patch [38]	12.7°	10.5°	11.3°	2.5°	26.2°
Grey-World [7]	13.0°	11.0°	11.5°	3.1°	26.0°
general Grey-World [3]	11.6°	9.7°	10.2°	3.4°	22.7°
1st – order Grey-Edge [48]	10.6°	8.8°	9.2°	3.0°	21.1°
Pixel-based Gamut Mapping [23]	11.8°	8.9°	10.0°	2.8°	24.9°
Edge-based Gamut Mapping [25]	12.8°	10.9°	11.4°	3.6°	25.0°
ML [8]	10.3°	8.9°	9.2°	2.8°	20.3°
Our method ($\delta = 0.0075$)	10.5°	8.2°	8.9°	2.2°	22.5°
Our method (fixed parameters)	11.1°	9.1°	9.5°	2.6°	23.0°

this matrix in our experiments to convert images into the sRGB colorspace before applying our algorithm. At this stage, it is important to note that the sensitivity functions of the camera used to create the database do not provide a perfect matching between the camera colorspace and sRGB. This is illustrated by the right part of Figure 3: when the groundtruth RGB triplets provided in the camera colorspace are converted into sRGB and projected in the xy diagram, the projections are significantly different from those obtained directly from the illuminant spectra. This is explained by the fact that the camera sensitivity functions are not fully compatible with the CIE standard observers color matching functions. Note that this is a drawback of our algorithm. Nonetheless, this problem should not occur with recent cameras, which now always offer sRGB compatibility.

3) *Comparative evaluation*: Tables I to V show the results obtained with the three different databases. In each table, the results of our approach are provided with two different sets of parameters: a set of fixed parameters, identical for all experiments and provided in Algorithm 1, and a set where the threshold δ is optimized for the database.

a) *Colorchecker*: For the Colorchecker dataset, results are provided both in the sRGB colorspace and in the colorspace of the camera (see Table I). Results from state of




Fig. 3. Groundtruth illuminants in the xy diagram for the SFU LAB database. On the left, the direct projection of illuminant spectra. On the right, projections of the illuminant triplets after conversion to sRGB. The difference between the projections is explained by an imperfect compatibility between the camera colorspace and sRGB.

the art approaches in the camera colorspace are reproduced from [26]. In order to provide the performances of our algorithm in this space, we project our estimations from sRGB to the camera RGB colorspace. As explained before, these results should be taken with caution, essentially because these errors are measured in two different camera colorspace, one corresponding to the Canon 5D, and the other to the Canon 1D. The RGB components of both cameras do not necessarily have the same range. As a consequence, averaging these errors is not necessarily consistent.

For this reason, we also provide the error statistics of all methods in sRGB. In order to compute these statistics, illuminant estimations are projected from the camera colorspace to sRGB for all state of the art algorithms. Observe that in both colorspace, our approach shows better results than standard, learning-free methods. We also see from this table that our approach provides comparable performances with more complex learning-based algorithms, while keeping a very reasonable computational cost.

b) Grey-Ball: On the Grey-Ball database (see Table II), our approach also shows better results than standard illuminant estimation methods. On this database, the results of our approach are even slightly better than the more complex, state of the art learning-based methods if the parameter δ is well chosen. Observe also that the algorithm is quite stable, providing good results both in the Colorchecker and in the Grey-Ball databases with the same set of parameters.

c) SFU Lab: For the SFU Lab database, results are provided both in sRGB and in the camera colorspace, in order to compare the method to state of the art approaches. Table III first shows the statistical errors on the entire database. As said before, these results should be interpreted with caution because the RGB color space of the camera is not perfectly compatible with sRGB, which is a prerequisite of our algorithm. Despite this drawback, the performance of our approach on this database is relatively close to those of the most recent state of the art algorithms. Looking at the median and trimean columns, we see that our approach yields slightly lower performances than the best learning-based approaches. On the Best 25% column, our results are slightly better. This can be explained by a large variance of the error distribution on this database. When a scene satisfies the algorithm assumptions, that is if there is a grey surface in the scene, or a set of very light grey dots, the algorithm yields excellent results. Otherwise, the results are very poor, as shown in the Worst-25% column. This is illustrated by Table V, which shows the errors statistics depending on the kind of observed scene. As might be expected, the scenes containing achromatic surfaces, shown on the top of Figure 4, produce excellent results. On the contrary, scenes without achromatic surfaces are subject to very large estimation errors. Table IV shows that error statistics are much less dependent on the type of illuminant used than on the type of scene.

C. Multi illuminant estimation

In this section, we evaluate the ability of our approach to deal with scenes lit by multiple light sources. As described

TABLE IV
PERFORMANCES ON THE SFU LAB DATABASE WITH RESPECT TO THE ILLUMINANT.

Illuminant	Mean	Median	Worst-25%(μ)
ph-ulm	8.5°	4.6°	21.5°
solux-3500+3202	5.3°	1.9°	16.6°
solux-3500	7.2°	2.7°	22.6°
solux-4100+3202	4.8°	1.6°	15.7°
solux-4100	6.5°	2.9°	20.3°
solux-4700+3202	4.5°	1.5°	15.2°
solux-4700	7.1°	2.7°	21.1°
syl-50MR16Q+3202	5.9°	2.1°	18.7°
syl-50MR16Q	7.3°	2.9°	22.4°
syl-cwf	6.9°	3.5°	20.7°
syl-wwf	5.9°	4.6°	14.6°

TABLE V
PERFORMANCES ON THE SFU LAB DATABASE WITH RESPECT TO THE TYPE OF SCENE.

Scene	Mean	Median	Worst-25%(μ)
Apples	21.4°	22.5°	29.7°
Blocks	27.4°	29.3°	31.2°
Fruit	20.0°	20.4°	25.7°
Paper	14.0°	12.5°	20.2°
Books	4.2°	3.3°	7.3°
Jersey	1.7°	1.6°	2.5°
Monkey	0.8°	0.8°	1.7°
Plastic	2.0°	1.0°	4.9°

in Section IV, the mode selection occurring at the end of our algorithm can be extended to detect several modes in the temperature histogram H . In practice, we assume that each of these detected modes has been generated by a light source.

Observe that removing the uniqueness assumption for the illuminant is potentially dangerous. Indeed, the inherent ambiguity of the illuminant estimation problem (an observed color may correspond to a neutral surface seen under a colored illuminant as well as a colored surface seen under a canonical illuminant) increases when the number of light sources is unknown.

Evaluating the algorithm in this multi-illuminant framework raises two questions. First, when only one light source is used in the scene, to what extent the multiple mode detection can deteriorate performances? Second, when several illuminants are indeed involved in the scene, how the results are improved by detecting these secondary modes? We suggest an experiment for each of these questions.

1) Measuring errors between two sets of illuminants: Before describing these experiments, let us detail how we measure the error between an arbitrary number of illuminants and estimations. For this task, we make use of the EMD distance introduced by Rubner *et al.* in [44]. Let N^g be the number of groundtruth illuminants present in the scene and let N^e be the number of estimated illuminants. We attribute a mass $1/N^g$ to each groundtruth illuminant and a mass $1/N^e$ to each one of our estimations. We then compute the EMD distance between the two sets of illuminants, the cost function between two illuminants being defined as their euclidean distance in the rg chromaticity diagram.

2) Multi-illuminant case: There is no standard dataset for the evaluation of multi-illuminant detection algorithms.



Fig. 4. Examples of scenes for which the algorithm hypotheses are satisfied (top) or not (bottom).

TABLE III
THE FIRST TABLE SHOWS PERFORMANCES ON THE SFU LAB DATABASE, IN THE CAMERA RGB COLORSPACE. THE SECOND TABLE SHOWS THE COMPARATIVE RESULTS IN THE UNIVERSAL SRGB COLORSPACE.

Method	Mean	Median	Trimean	Best-25%	Worst-25%
White-Patch [38]	9.1°	6.5°	7.5°	1.8°	21.1°
Grey-World [7]	9.8°	7.0°	7.6°	0.9°	23.6°
general Grey-World [3]	5.4°	3.3°	3.8°	0.5°	13.8°
1st – order Grey-Edge [48]	5.6°	3.2°	3.7°	1.1°	14.1°
Pixel-based Gamut Mapping [23]	3.7°	2.3°	2.5°	0.5°	9.4°
Edge-based Gamut Mapping [25]	3.9°	2.3°	2.7°	0.5°	10.0°
Our method ($\delta = 0.015$)	6.4°	2.4°	3.6°	0.3°	19.3°
Our method (Fixed parameters)	6.5°	2.7°	4.0°	0.3°	19.0°

Method	Mean	Median	Trimean	Best-25%	Worst-25%
White-Patch [38]	8.0°	5.6°	6.3°	1.8°	18.1°
Grey-World [7]	10.8°	7.8°	8.3°	1.2°	26.6°
general Grey-World [3]	5.0°	2.7°	3.5°	0.6°	13.0°
1st – order Grey-Edge [48]	5.4°	3.1°	3.7°	1.0°	13.5°
Pixel-based Gamut Mapping [23]	4.2°	2.5°	2.9°	0.6°	10.6°
Edge-based Gamut Mapping [25]	4.5°	2.9°	3.3°	0.7°	11.4°
Our method ($\delta = 0.015$)	5.3°	1.9°	2.5°	0.3°	16.4°
Our method (Fixed parameters)	6.2°	2.6°	3.7°	0.4°	17.7°

In order to evaluate our algorithm, we follow the protocol previously proposed by Gijsenij *et al.* in [29]. We create a synthetic dataset from the Colorchecker database. Each image is first corrected using the provided groundtruth illuminant and the Von Kries model. Then, for each corrected image, two illuminants L^1 and L^2 are randomly selected¹⁰ in the set of groundtruth illuminants provided with the database. The corrected image is then illuminated by these two light sources, a mask defining the local proportions of each illuminant, as illustrated by Figure 5. Several masks are possible (vertical, horizontal or diagonal). In order to smooth the transition between illuminants, Gaussian filtering with random standard deviation (between 1 and 5) is used to blur the binary masks.

3) *Single illuminant case:* For the single illuminant case, we reproduce partially the experiment on the Colorchecker

¹⁰We check that these two illuminants are not too close in practice, and discard the pairs such that $\text{EMD}(L^1, L^2) > 1.11$. This threshold roughly corresponds to the one used in [29].

dataset [47] described in Section V-B, except that we allow our algorithm to detect several meaningful modes. The experiment illustrates the behavior of our algorithm when the *a priori* on the number of illuminant is removed. The performance of the algorithm can only decrease in this framework, the goal of the experiment is to measure to which extent.

4) *Results:* Table VI shows the performance of our algorithm on the set of images generated using two light sources, as described in V-C2. Our algorithm is tested twice on this dataset. In the first experiment, we deliberately keep only one mode (the most meaningful one) in the histogram. In the second experiment, we release the uniqueness assumption and we keep all the maximal meaningful modes. In this multi-illuminant framework, authorizing the detection of multiple modes permits to increase quite significantly the detection accuracy.

Table VII shows the performance loss resulting from the relaxation of the uniqueness assumption when applying our




Fig. 5. Left: example of an image created for estimating multiples light sources. Right: mask used to define the spatial proportion of each illuminant.

algorithm on the original images of the Colorchecker dataset. The first line of the table shows the error statistics when the number of illuminants to be detected is not imposed at all, and the second line shows the performance of the single mode detection. Observe that the performance loss between both configurations is very reasonable. We insist on the fact that to the best of our knowledge, no algorithm had been proposed in the literature to estimate at the same time the number of illuminant present in the scene and the color of these illuminants.

VI. CONCLUSION

In this paper, a new procedure for the estimation of single or multiple illuminants in color images has been proposed. A novel and simple voting scheme has been introduced and combined with physical constraints, in order to select the most probable illuminants in a chromaticity space. It has been shown that this algorithm permits to achieve state of the art performances on different challenging databases without requiring any learning step. The ability of the whole procedure to deal with situations where the number of light sources is unknown has also been demonstrated. To the best of our knowledge, such ability was not reported before in the literature. The proposed procedure works globally on images, even in the multi-illuminant case. This strongly differs from several recent approaches working locally on image patches or superpixels [5], [29]. These methods enable local color corrections, but cannot automatically detect the number of light sources. As a future work, we intend to explore connexions and tradeoffs between these local procedures and our approach.

REFERENCES

- [1] Kobus Barnard, Vlad Cardei, and Brian Funt. A comparison of computational color constancy algorithms. i: Methodology and experiments with synthesized data. *Image Processing, IEEE Transactions on*, 11(9):972–984, 2002.
- [2] Kobus Barnard, Graham Finlayson, and Brian Funt. Color constancy for scenes with varying illumination. *Computer vision and image understanding*, 65(2):311–321, 1997.
- [3] Kobus Barnard, Lindsay Martin, Adam Coath, and Brian Funt. A comparison of computational color constancy algorithms. *IEEE Trans. Image Process.*, 11:2002, 2002.
- [4] Simone Bianco, Gianluigi Ciocca, Claudio Cusano, and Raimondo Schettini. Improving color constancy using indoor–outdoor image classification. *Image Processing, IEEE Transactions on*, 17(12):2381–2392, 2008.
- [5] Michael Bleier, Christian Riess, Shida Beigpour, Eva Eibenberger, Elli Angelopoulou, Tobias Trger, and Andr Kaup. Color constancy and non-uniform illumination: Can existing algorithms work? In *ICCV Workshops*, 2011.
- [6] Ivaylo Boyadzhiev, Kavita Bala, Sylvain Paris, and Frédo Durand. User-guided white balance for mixed lighting conditions. *ACM Trans. Graph.*, 31(6):200:1–200:10, November 2012.
- [7] G Buchsbaum. A spatial processor model for object colour perception. *J. Franklin Inst.*, 310(1):1–26, 1980.
- [8] Ayan Chakrabarti, Keigo Hirakawa, and Todd Zickler. Color constancy with spatio-spectral statistics. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2012.
- [9] Florian Ciurea and Brian Funt. A large image database for color constancy research, 2003.
- [10] Dave Coffin. Decoding raw digital photos in linux, 2013.
- [11] A. Desolneux, L. Moisan, and J.-M. Morel. A grouping principle and four applications. *PAMI*, 25(4):508–513, 2003.
- [12] A. Desolneux, L. Moisan, and J. M. Morel. *From Gestalt Theory to Image Analysis*, volume 34. Springer-Verlag, 2008.
- [13] Marc Ebner. A parallel algorithm for color constancy. *J. Parallel Distrib. Comput.*, 64(1):79–88, January 2004.
- [14] Marc Ebner. *Color constancy*, volume 6. Wiley, 2007.
- [15] Marc Ebner. Estimating the color of the illuminant using anisotropic diffusion. In Walter G. Kropatsch, Martin Kampel, and Allan Hanbury, editors, *Computer Analysis of Images and Patterns, 12th International Conference, CAIP 2007, Vienna, Austria, August 27-29, 2007, Proceedings*, volume 4673 of *Lecture Notes in Computer Science*, pages 441–449. Springer, 2007.
- [16] G. D. Finlayson and S. D. Hordley. Gamut constrained illuminant estimation. *Int. J. Comput. Vision*, 67:2006, 2006.
- [17] Graham D. Finlayson, Mark S. Drew, and Brian V. Funt. Color constancy: Generalized diagonal transforms suffice. *J. Opt. Soc. Am. A*, 11:3011–3020, 1994.
- [18] Graham D. Finlayson, Brian V. Funt, and Kobus Barnard. Color constancy under varying illumination. In *ICCV*, pages 720–725, 1995.
- [19] Graham D. Finlayson and Steven D. Hordley. Color constancy at a pixel. *J. Opt. Soc. Am. A*, 18(2):253–264, Feb 2001.
- [20] Graham D. Finlayson, Steven D. Hordley, and Paul M. Hubel. Color by correlation: A simple, unifying framework for color constancy. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 23(11):1209–1221, 2001.
- [21] Graham D. Finlayson and Gerald Schaefer. Solving for colour constancy using a constrained dichromatic reflection model. *Int. J. Comput. Vision*, 42(3):127–144, May 2001.
- [22] Graham D. Finlayson and Elisabetta Trezzi. Shades of gray and colour constancy. In *Color Imaging Conference*, pages 37–41, 2004.
- [23] D. A. Forsyth. A novel algorithm for color constancy. *Int. J. Comput. Vision*, 5(1):5–36, 1990.
- [24] P. V. Gehler, C. Rother, A. Blake, T. Minka, and T. Sharp. Bayesian color constancy revisited. In *Conference on Computer Vision and Pattern Recognition*, pages 1–8, 06 2008.
- [25] A. Gijsenij, T. Gevers, and J. van de Weijer. Generalized gamut mapping using image derivative structures for color constancy. *Int. J. Comput. Vision*, 86(2-3):127–139, 2010.
- [26] A. Gijsenij, T. Gevers, and J. van de Weijer. Computational color constancy: Survey and experiments. *IEEE Trans. Image Process.*, 20(9):2475–2489, 2011.

TABLE VI

RESULTS ON IMAGES GENERATED USING TWO LIGHT SOURCES (SEE SECTION V-C2). FIRST LINE: THE ALGORITHM ASSUMES ONE SINGLE LIGHT SOURCE. SECOND LINE: THE UNIQUENESS ASSUMPTION IS REMOVED. $k = 3$, $\delta = 0.0150$, $T_m = 2000$, $T_M = 20000$, $bins = 300$

Method	Crit.	Mean	Med	Trimean	Best-25%	Worst-25%	Max.
Single mode	5.8	6.2	5.3	5.6	1.3	12.6	26.1
Multiple modes	3.7	4.5	2.9	3.4	0.3	11.0	20.3

TABLE VII

RESULTS ON THE COLORCHECKER DATASET WHEN NO PRIOR INFORMATION ON THE NUMBER OF ILLUMINANTS IS PROVIDED. $k = 3$, $\delta = 0.0150$, $T_m = 2000$, $T_M = 20000$, $bins = 100$

Method	Crit.	Mean	Med	Trimean	Best-25%	Worst-25%	Max.
Multiple modes	3.9	4.7	3.1	3.6	0.6	11.2	23.1
Single mode	2.8	3.7	2.0	2.4	0.5	9.6	25.6

- [27] A. Gijsenij, Th. Gevers, and M.P. Lucassen. A perceptual analysis of distance measures for color constancy. *Journal of the Optical Society of America A*, 26(10):2243–2256, 2009.
- [28] A. Gijsenij, Th. Gevers, and J. van de Weijer. Physics-based edge evaluation for improved color constancy. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition 2009 (CVPR'09)*, pages 1–8, Miami Beach, Florida, USA, June 2009.
- [29] A. Gijsenij, Rui Lu, and T. Gevers. Color constancy for multiple light sources. *Image Processing, IEEE Transactions on*, 21(2):697–707, feb. 2012.
- [30] Arjan Gijsenij and Theo Gevers. Color constancy using natural image statistics and scene semantics. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(4):687–698, 2011.
- [31] Javier Hernández-Andrés, Raymond L. Lee, and Javier Romero. Calculating correlated color temperatures across the entire gamut of daylight and skylight chromaticities. *Appl. Opt.*, 38(27):5703–5709, Sep 1999.
- [32] Eugene Hsu, Tom Mertens, Sylvain Paris, Shai Avidan, and Frédo Durand. Light mixture estimation for spatially varying white balance. *ACM Trans. Graph.*, 27(3):70:1–70:7, August 2008.
- [33] Yoshie Imai, Yu Kato, Hideki Kadoi, Takahiko Horiuchi, and Shoji Tomimaga. Estimation of multiple illuminants based on specular highlight detection. In *Proceedings of the Third international conference on Computational color imaging, CCIW'11*, pages 85–98, Berlin, Heidelberg, 2011. Springer-Verlag.
- [34] Deane B. Judd, David L. Macadam, Günter Wyszecki, H. W. Budde, H. R. Condit, S. T. Henderson, and J. L. Simonds. Spectral distribution of typical daylight as a function of correlated color temperature. *J. Opt. Soc. Am.*, 54(8):1031–1036, Aug 1964.
- [35] Rei Kawakami, Katsushi Ikeuchi, and Robby T. Tan. Consistent surface color for texturing large objects in outdoor scenes. In *Proceedings of the Tenth IEEE International Conference on Computer Vision - Volume 2, ICCV '05*, pages 1200–1207, Washington, DC, USA, 2005. IEEE Computer Society.
- [36] Rei Kawakami, Robby T Tan, and Katsushi Ikeuchi. A robust framework to estimate surface color from changing illumination. In *Asian Conference on Computer Vision (ACCV2004)*, 2004.
- [37] Hyuk-Ju Kwon, Sung-Hak Lee, Tae-Wuk Bae, and Kyu-Ik Sohng. Compensation of de-saturation effect in hdr imaging using a real scene adaptation model. *J. Visual Commun. Image Represent.*, (0), 2012.
- [38] E.H. Land and John J. McCann. Lightness and retinex theory. *J. Opt. Soc. Am.*, 61(1):1–11, Jan 1971.
- [39] Bing Li, De Xu, Weihua Xiong, and Songhe Feng. Color constancy using achromatic surface. *Color Research & Application*, 35(4):304–312, 2010.
- [40] Baptiste Mazin, Julie Delon, and Yann Gousseau. Illuminant estimation from projections on the planckian locus. In *ECCV Workshops (2)*, pages 370–379, 2012.
- [41] Irwin G. Priest. A proposed scale for use in specifying the chromaticity of incandescent illuminants and various phases of daylight. *J. Opt. Soc. Am.*, 23(2):41–45, Feb 1933.
- [42] Whitman Richards and Edward A Parks. Model for color conversion. *JOSA*, 61(7):971–976, 1971.
- [43] Charles Rosenberg, Martial Hebert, and Sebastian Thrun. Color constancy using kl-divergence. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, volume 1, pages 239–246. IEEE, 2001.
- [44] Y. Rubner, C. Tomasi, and L. J. Guibas. A metric for distributions with applications to image databases. In *Proceedings of the 1998 IEEE International Conference on Computer Vision*, pages 59–66, January 1998.
- [45] Guillermo Sapiro. Color and illuminant voting. *IEEE Trans. Pattern Anal. Mach. Intell.*, 21:1210–1215, 1999.
- [46] Steven A. Shafer. Color. In Glenn E. Healey, Steven A. Shafer, and Lawrence B. Wolff, editors, *Using color to separate reflection components*, pages 43–51. Jones and Bartlett Publishers, Inc., USA, 1992.
- [47] Lilong Shi and Brian. Re-processed version of the gehler color constancy dataset of 568 images, 2010.
- [48] J. Van De Weijer, T. Gevers, and A. Gijsenij. Edge-based color constancy. *IEEE Trans. Image Process.*, 16(9):2207–2214, 2007.
- [49] Joost Van De Weijer, Cordelia Schmid, and Jakob Verbeek. Using high-level visual information for color constancy. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pages 1–8. IEEE, 2007.
- [50] Javier Vazquez-Corral, Maria Vanrell, Ramon Baldrich, and Francesc Tous. Color constancy by category correlation. *IEEE Transactions on Image Processing*, 21(4):1997–2007, April 2012.
- [51] Günther Wyszecki and W. S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulae*. Wiley-Interscience, 2 edition, August 2000.